

INTEL CORPORATION
5200 N.E. Elam Young Parkway
Hillsboro, Oregon 97124-9987
(503) 681-8080



Dear Customer:

Thank you for purchasing the Developer Version of Intel UNIX Release 4.0. The Developer Version is functionally complete but is not bug-free. This product is intended for use primarily in developing and porting application or device driver software.

Please be sure to register this product with Intel so you will be eligible to purchase the Production Version of Intel UNIX 4.0 at a substantial discount.

Intel is pleased to be selling and supporting the industry standard UNIX for 386 and i486 based computers. Intel's UNIX Release 4.0 brings together System V, Xenix, and BSD systems.

Application binaries that run on UNIX Release 3.2 also run on UNIX Release 4.0. Most Xenix applications also run on Intel's Release 4.0. UNIX Release 4.0 from Intel provides many BSD/SunOS features and delivers a high degree of BSD/SunOS source level compatibility.

The product you've received includes a complete two user UNIX system, ILS support, Installation Guide, Release Notes, Product Overview and Master Index, and the following software:

- * ST506/ESDI Boot Disk (5 1/4 and 3 1/2)
- * WD SCSI Boot Disk (5 1/4 and 3 1/2)
- * Adaptec SCSI Boot Disk (5 1/4 and 3 1/2)
- * Base System Disk 2 (5 1/4 and 3 1/2)
- * Hardware Driver Set (5 1/4 and 3 1/2)
- * UNIX 4.0 Installation Tape (60 MB)

Document Number
466226

Sheet Rev
1 A



Please call the appropriate number from the list below if any of these items are missing.

U.S.	800-INTEL4U
U.K.	(0) 793 / 641469
France	(16) 1 / 30.57.70.00
Germany	(0) 89 / 903-2025
Netherlands	(0) 10 / 4.07.11.30
Italy	(0) 2 / 892.00950
Israel	(0) 3 / 548-3222
Sweden	(0) 8 / 825416
Finland	(9) 0 / 544644
Spain	(9) 1 / 308-2552
Switzerland	call Germany
Denmark	call Sweden
Norway	call Sweden
Other countries	call nearest facility

A software product is only as good as the support behind it, especially a product as powerful and complex as the UNIX operating system. We have included our Installation Level Support (ILS) at no charge with your purchase. ILS is your assurance that you will be able to successfully install your Intel UNIX software.

ILS includes replacement of defective media within 90 days of purchase, one hour of installation support by phone, and access to Intel's UNIX bulletin board.

ILS also includes installation support for these related Intel UNIX software products:

Product Code	Description
UNIX40ULTDLC	Multiuser Upgrade
UNIX40OLDOC	On-line Documentation
MERGE40	DOS under UNIX capability
MOTIF40	Motif Graphical Interface

With ILS, support engineers will help you to get a shell prompt on the system console. ILS does not include network connection, application configuration, system administration, or application support.

Please call the appropriate number listed above to receive phone installation support. Hours of operation for U.S. and Canada are 8:00am to 5:00pm PST. Other locations



operated during normal business hours. All local holidays are observed.

Write down your system's hardware configuration before you call - including attached hardware devices and add-in cards (manufacturer, model, and revision level are all helpful to the support engineer).

Fill out the product registration card, and mail it using the address label for your nearest Intel office.

By registering your products, you will receive bulletin board access, product update notices, and UNIX training information. Dial the bulletin board in the U.S. at (503) 640-3042. The bulletin board supports modems up to 2400 baud; you'll find the latest product information and various articles on the bulletin board.

Your initial session on the bulletin board is restricted to general public access, but you can request more access privileges. Log into the bulletin board, and choose the request menu to request upgraded privileges. Your access privileges will be upgraded by the next business day and will expire 30 days later.

Contact your Intel salesperson or Intel distributor for information on other Intel UNIX software products.

Thanks for choosing Intel.





*product
release
notes*

**Intel UNIX System V
Release 4.0**

1. Installation Notes

This section contains information needed before installing Intel UNIX¹ System V Release 4.0 (SVR4) on your system. The full installation procedure is documented in the *System V Release 4.0 Installation Guide*.

1.1 Installation

Be sure to change directory to root (/) prior to installing any package. Unexpected results may occur if installing from somewhere other than root.

1.2 Boot Disks

Boot disks are now write-protected (they were not in previous products). This is normal and the write-protect tabs should not be removed.

1.3 Overlay Installation

Intel UNIX System V Release 4.0 cannot be installed as an "upgrade" or overlay on top of previous versions of UNIX System V, including pre-release versions of SVR4. For complete installation instructions, please refer to the *Installation Guide*. The overlay procedure described in the *Installation Guide* is only supported for re-installing the same version of SVR4.

1.4 Memory greater than 16 megabytes

Memory configurations above 16 megabytes are not supported in this release.

1.5 Size of Swap

When running networking applications, the size of /dev/swap should be at least two times your system memory. You can allocate swap space during installation. After installation, use the swap -a command to add additional space.

1.5.1 Swap and System Hang

If you use the -a option of the swap command to add a regular file as a swap area (see swap(1M), your operating system may hang if it runs out of memory while trying to manipulate swap space in the file. It is recommended that you not use a regular file as a swap area.

1. UNIX is a registered trademark of AT&T.

1.6 Configuring an s52k Filesystem

If you configure your root file system to of type s52k, the following warning message will be displayed each time the system is booted:

WARNING: nfs_mountroot called

This warning can be ignored. You can stop the message from being displayed by changing a line in the file `/stand/boot`. Change the line that reads:

`rootfstype=s52k`

to read:

`rootfstype=s5`

1.7 Non-SCSI Hardware Controller Configuration

Before installing Intel UNIX System V Release 4.0, verify that the controllers for ESDI machines and the host adapter for SCSI machines are configured properly according to the documentation that came with your controller.

1.7.1 ESDI Disk Controllers

A common use of ESDI disk controllers is to run them with the on-board BIOS removed. Running in this manner, System V Release 4.0 can't recognize cylinders above 1024, thus making large disks appear smaller than they are. If the correct system set up value does not recognize the entire disk, replacing the BIOS will solve this problem.

1.7.2 Supported Cartridge Tape Controllers

The current set of tape controllers and tape drives supported by the cartridge tape driver package (`qt`) includes:

Tape Controller	Supported Tape Drive	Comments
Wangtek PC-36	Wangtek 5099-EN	Standard, 60MB QIC-24
Everex 811 (Bell Tech)	Wangtek 5150-EN	QIC-24/120/150
Everex 811 (Bell Tech)	Archive Viper 2150L	QIC-24/120/150
Archive SC499-R	Archive Ext. FT-60	60MB QIC-24
Archive VP402	Archive Viper 2150L	QIC-24/120/150
Archive VP402	Wangtek 5150-EN	QIC-24/120/150

The correct settings for the base controller address are 0x300, interrupt vector 5 and DMA channel 1.

If you are using Archive controllers and Archive tape drives or an Everex controller/Archive tape drive combination, you must use the cartridge tape driver package (qt).

If your controller is not configured correctly, the following message is displayed:

*Cartridge Controller was not found at address 000000300H
Tape driver disabled*

The cartridge tape will be inoperative.

1.8 SCSI Configuration Considerations

If you're running a SCSI-based system with a SCSI tape drive, do not install a non-SCSI tape controller and the qt cartridge tape software. They both try to use the same nodes in /dev/rmt.

The following is not supported:

1. Two Adaptec SCSI controllers in the same system.
2. ESDI or ST506 (non-SCSI) with either a Western Digital or Adaptec (SCSI) combination.

1.9 Configuring Two WD-7000 FASST SCSI Host Adapter Cards

Two WD-7000 FASST² SCSI Host Adapter cards can be installed into your system. Three cards should also work, but four cards cannot be supported since the only DMA channels the hardware and the software can use are channels 5, 6 & 7.

For experienced system administrators, here's how to configure two WD7000s (CAUTION: this procedure requires reconfiguration of the SCSI controller board):

1. Use default configuration for the first WD7000:

- IRQ 15
- DACK/DRQ 6
- BIOS ROM address 0xCE000
- I/O address range 350-353

2. FASST and Western Digital are trademarks of Western Digital Corporation

2. Use following configuration for the second WD7000:

- IRQ 11
- DACK/DRQ 5
- BIOS ROM address 0xC8000
- I/O address range 330-333
- Edit the `/etc/conf/sdevice.d/scsi` file and add the lines:

SCSI	Y	1	5	1	15	350	353	0	0
SCSI	Y	1	5	1	11	330	333	0	0

- Rebuild the UNIX kernel by executing `/etc/conf/bin/idbbuild`.
- Remove the BIOS ROM from the second WD7000 (U60 near jumper W7).
- Configure the jumpers on the second WD7000 as:

W6: in and W9 in to disable the floppy hardware,
W1: out out out out, as before, for IRQ Channel,
W2: out out in out out in out out in out, for IRQ 11,
W3: in out out in in, for I/O address range 0x330-0x333,
W4: in in out in, for BIOS ROM address 0xC8000 (not used).

- Install both WD7000s in 16 bit expansion slots.
- Connect a SCSI cable to each WD7000.
- Connect other SCSI devices, make sure that each SCSI cable is properly terminated at the ends.
- Reboot.

You may need to manually remake the device nodes depending on the devices hooked up to the two SCSI buses. For example, if you have on each SCSI bus a WD7000 (ID 7), a hard disk (ID 0), and a tape (ID 3):

<code>/dev/dsk/c0t0d0s0</code>	as always is your system disk on BUS-0
<code>/dev/rmt/c0s0</code>	is the tape drive on BUS-0
<code>/dev/dsk/c0t1d0s0 or ls0</code>	happens to be the disk on BUS-1
<code>/dev/rmt/c0s1</code>	is the tape drive on BUS-1

The last two should be /dev/dsk/c1t0d0s0 and /dev/rmt/c1s0. To change, edit the /etc/conf/node.d files.

If you have three disks or three tapes in a system, you need to make the nodes for the third devices manually. The minor number for each subsequent device is always 16 greater than the previous one. So, the minor number for the 3rd device is 16 greater than the minor number for the 2nd device.

1.10 Supported Extended VGA Video Cards

The following popular video cards are known to support extended VGA (EVGA) mode (i.e., 800x600 pixels):

- Orchid Pro Designer VGA
- Paradise VGA Plus 16
- STB VGA Extra/EM 16
- Video7 VRAM VGA

1.11 Reinstallation of SCDE

If the scde (Standard C Development Environment) package is removed from the system and you wish to reinstall the package, /etc/profile must be saved before you reinstall. The reinstallation will indicate a partial failure, which you may ignore. After the reinstallation is done, copy back the saved /etc/profile to /etc/profile.

1.12 Dynamic Libraries

Most of the executables in Intel UNIX System V Release 4.0 are compiled with dynamic libraries. To determine which dynamic library (if any) the executable uses type in:

```
mcs -p -n .interp executable
```

This will display the dynamic library being used with this executable. It will display nothing if the executable is compiled with archive (ar) libraries.

1.13 Network Controller Configuration

See the section *Configuring Multiple Ethernet Boards* in the Networking Notes sections of this document.

2. Operating System Notes

2.1 Booting with initdefault level 1

When booting using an `initdefault` level of 1, some of the start/stop scripts may fail to detect that the previous run level was 0, and attempt to run commands for shutting down, resulting in error messages that you may ignore.

2.2 File-System-Independent Booting

The BFS file system type does not support the following features (`/stand` is a BFS file system):

- `mmap`
- `truncate up`
- symbolic links
- `volcopy`
- `labelit`
- `mkdir`

2.3 Init behavior when executed inside a vt

Do not attempt to execute the `init` command from a `vt` shell other than the console (`/dev/console` or `/dev/vt00`).

2.4 Debugging and vt Devices

If the kernel debugger, `kdb`, is entered manually (CTL-ALT-D) while the active terminal is a `vt` (i.e. not `/dev/console`), the debugger is started on the console and is not switched from the `vt` session to the console, leading to the appearance of a hung system. Entering a `q` and <return> quits the kernel debugger and resume operation of the system. Therefore, do not start the kernel debugger while running on devices other than `/dev/console`. Note that the kernel debugger (`kdb`) is an optional package.

2.5 Simultaneous Cartridge Tape and Floppy use

Systems may experience floppy read/write failures when the floppy access occurs during a data transfer to/from an integral non-SCSI 1/4" cartridge tape unit, or heavy display activity.

2.6 Device Driver Specification

The DDI/DKI (Device Driver Interface/Driver Kernel Interface) Reference Manual describes access functions for certain kernel data structures. It also lists include

header files that are permitted. From the list of include files permitted, one might conclude that it was acceptable to reference all the global data structures declared in those include files. Doing so may adversely restrict potential future implementations of the kernel. Therefore, it is recommended that no global data be used.

2.7 bp_mapin returns void

The manual page documents `bp_mapin(D3DK)` routine as returning a type `vaddr_t`. It is really a function returning type `void`.

2.8 Mouse Driver Support of Multiple Terminal Assignments

A compilation problem is preventing the mouse driver from supporting a configuration where one mouse is mapped to the console and another to a second console-like device. If you need to assign a mouse to this second console device, deconfigure the mouse from the console.

2.9 ioctl Command Conflicts

There are conflicts between `ioctl` command values in `kd.h`, `streamio.h`, and `termiox.h`.

The commands `SPECIAL_IOPRIVL`, `STSET`, and `TCGETX` all have the same value. This doesn't cause a problem only because there are no drivers that recognize `STSET` and `TCGETX`, and support for `SPECIAL_IOPRIVL` is not yet implemented. Once support is added for `SPECIAL_IOPRIVL` and/or drivers are added that recognize `STSET` or `TCGETX`, these commands will not be handled correctly.

There are also potential conflicts between the following pairs of `ioctl` commands:

- `STTHROW` and `TCSETEX`
- `STWLINE` and `TCSETEXW`
- `STTSV` and `TCSETEXF`.

The `ioctl` command values in `streamio.h` and `termiox.h` will be changed in the next release to eliminate these conflicts. This will require a recompilation.

2.10 Utilities and SCSI Drivers

Not all utilities work properly with SCSI disk drivers. For example, the `sar` (system activity report) and `sadp` (disk access profiler) utilities do not support the SCSI disk driver (`sd01`).

2.11 Cannot Cross Layers in shl when -loblk is off

Do not attempt to display the output of a layer in another layer in shell layers when the `stty` option `-loblk` is off.

2.12 uucp and uuto Limitations

`uucp` and `uuto` cannot properly transfer any file larger than 2 Megabytes (4096 blocks). If you need to transfer a file larger than 4096 blocks, use `dd` with the `skip` and `count` options to split the file up. Example:

```
dd if=big.file of=tmp.0 count=4096  
dd if=big.file of=tmp.1 count=4096 skip=4096  
dd if=big.file of=tmp.2 count=4096 skip=8192  
.  
. .  
uuto tmp.* machx!user
```

2.13 /usr/ucb/mt Command

The `/usr/ucb/mt` command does not work and will not be supported in the future. Use `tapecntl` instead.

3. File Systems Notes

3.1 File System Creation Limitation

If you are creating file systems on a large (300 MB or larger) disk either at installation time or later, make sure that no file system takes up more than half the disk. Specifically, if either a `ufs` file system or an `s5` file system contains more than 65534 ($2^{16}-2$) inodes, various problems will result, in particular UNIX System V Release 3.2 binaries will not be able to `stat()` files with inodes that are greater than 65534.

3.2 Virtual File System

3.2.1 Future Directions

Aspects of the VFS interface (particularly the interaction between file system code and the Virtual Memory system) are subject to change in future releases of the system. Absolute compatibility is not guaranteed.

3.2.2 Switchout

The following commands were coded to work on the VFS switchout mechanism: `clri`, `dcopy`, `df`, `ff`, `ncheck`, `fsck`, `fsdb`, `labelit`, `mkfs`, `mount`, `umount`, and `volcopy`. These commands have a generic module and file system specific modules that the generic module calls. Commands operating on unmounted file systems either have the file system type supplied on the command line with the `-F` option, or an appropriate entry will be provided by the administrator in `/etc/vfstab` (see the *System Administrator's Guide* for more details).

3.3 File System Commands

3.3.1 clri Command

The `clri` command is currently working under the switchout mechanism (see *Switchout*) but will be dropped in a future release. Its functionality is being replaced by `fsdb -z`.

3.3.2 Future Direction of the mount Command

The format of the output of the `mount` command will be replaced in a future release by a new format (currently provided under the `-v` option).

Since the `ff` command is more flexible, it is intended that one day it will supersede `ncheck`. In a future release, `ff` will subsume all functionality of `ncheck` and `ncheck` will be removed from the system.

3.4 File System Types

3.4.1 /proc and /dev/fd

Even though they are not explicitly installed or discussed in the installation process, the `/proc` and `/dev/fd` file systems are created on every SVR4 system. This is normal, as is the fact that `dfspace` reports at login time that each has 0.00% (0.00 of 0.00 MB) available.

3.4.2 RFS

See Section 4, *Networking Notes*, for information on Remote File Sharing (RFS) implementation issues.

3.4.3 /proc File System

Utilities that use `/proc` can become hung if they try to read information about an existing hung process. They may sleep waiting to lock a process' information. If that

process happens to be hung, the utility may also sleep forever. This problem should only become apparent if a user executes a utility that does locking in /proc (such as ps or priocntl), and the system currently has a live process that is hung after locking its process information. The utility will continue to accept signals.

Reading of /proc files returns EIO error. If you are doing a cpio from /, unmount /proc first.

3.4.3.1 Known Problems

Given a corrupted UFS file system, if the file system is mounted read-only, fsck is executed, and then the remount option is used to mount the file system read-write, any changes fsck may have made will be destroyed (i.e. the file system will still be corrupted). This is because the remount option is handled differently for s5 than it is for UFS. It is expected that the behavior of UFS will be changed to match the behavior of s5 in this instance in a future release.

3.4.3.2 ufsdump/ufsrestore

ufsdump sometimes reports more blocks written than it actually writes. Therefore, the blocks reported by ufsrestore and ufsdump may be slightly different. The dump is executed correctly and should be usable. If there is a discrepancy between the two commands on the number of blocks in the dump, the number reported by ufsrestore, which should be slightly less than that reported by ufsdump, is correct.

3.4.3.3 mkfs and UFS

mkfs on a UFS file system has a compatibility mode that lets you limit the number of inodes to 64k. This provides compatibility with System V Release 3.2 applications that require access to inode numbers. This compatibility mode is used by default in disksetup.

3.4.3.4 quota Command

The quota command will not work correctly in all cases unless all UFS file systems that are currently mounted have a quota file. An empty file named quota should be placed at the root of each UFS file system.

3.5 Symbolic Links

The following list references a few of the incompatibilities between SVR4 and SunOS 4.0.

3.5.1 Scripts and Relative Symbolic Links

When a script is executed implicitly (e.g. with a `#!` as the first line in the script), and the command to be executed is a relative symbolic link to a relative symbolic link to the real file, then the script will not be found. The work around is to reference absolute symbolic links or hard links in scripts.

When a shell script is executed implicitly, with a `#!` as the first line in the script via a symbolic link, it will set `$0` to the name of the actual script rather than the name of the symbolic link. This may cause problems with scripts that check to see how they are called to determine their behavior. Currently, there is no workaround for this.

3.6 /etc/vfstab Entries Created by diskadd

If the `diskadd` command is run more than once, the entries in `/etc/vfstab` for that drive may not be deleted. If you are going to run the `diskadd` command more than once for a given drive, delete the entries for it in `/etc/vfstab` prior to running the `diskadd` command.

4. Networking Notes

4.1 TCP/IP (DARPA Internet Protocol Suite)

4.1.1 Domain Name Service (DNS)

The default Internet services, including `ftp` and `telnet`, are not compiled to use the Domain Name Service. For DNS, add `/usr/lib/resolv.so` to the `/etc/netconfig` file. The following is an example:

```
udp tpi_clts v inet udp /dev/ucp /usr/lib/tcpip.so,/usr/lib/resolv.so
```

4.1.2 netstat Command

SVR4 does not support the `-h` and `-t` options to the `netstat` command, as documented on the `netstat(1M)` manual page.

4.1.3 Pseudo Terminals

Pseudo terminals must be configured in order for TCP/IP services, such as `telnet` and `rlogin`, to work. These may be selected when the NSU package is installed.

4.2 NFS

4.2.1 Automounter

Direct map automounters do not always clean up after themselves. They may leave files that are symbolic links to automount daemon mount points for which the automount daemon no longer exists. Attempts to access these files will result in "*server not responding*" messages. These files can only be removed by unlinking them, then running `fsck`, or by rebooting and then removing any remaining links. The only other workaround is to avoid using direct maps with the automounter.

Note that the automount daemon for direct maps mounts itself on the mountpoint specified in the map, while the automounter for an indirect map mounts itself on the directory specified as the root for the indirect mount, i.e., one level above the mount point for the resource in the indirect map. This can be misleading to the user because the user may not expect anything to be mounted above the mount point of the resource. Because the daemon mounts itself there, the previous contents of the directory are covered for the duration of the life of the automounter. As with all file system mount points, it is a good policy to use empty directories as automount root mount points.

There is no entry in the `/etc/mnttab` file corresponding to the mount of the automount daemon on the automount root mount point. Therefore, it is possible for a user attempting to mount a file system on a given mount point to get a "*mount point busy*" message even though that directory is not noted in the `/etc/mnttab` file as an existing mount point.

4.2.2 Lock Manager

The lock manager does not always detect deadlock.

4.2.3 Loopback

If a large read or write operation to a loopback mount with `biod`'s running is interrupted, some of the `nfsd`'s can hang while waiting for the STREAMS code to free a STREAMS message block.

4.2.4 Secure NFS mount

If a secure NFS mount is interrupted while waiting for the keyserver to respond, it is possible for the file system to be mounted, but without a corresponding entry put in the `/etc/mnttab` file. This occurs because the signal will not be handled until after the mount system call has succeeded and returned to the `mount` command.

The command process will then be killed before it writes to /etc/mnttab.

4.2.5 Transport Independence

NFS can run over multiple transports simultaneously, provided that:

- a. both the client and the server agree on the maximum packet size for a given transport; or, the -o rsize=,wsize= options are given to the mount command to reduce the read and write sizes to a value small enough for both machines (the minimum of the packet sizes of the two machines' transports). Note that the read and write transfer sizes should be about 430 bytes smaller than the transport's packet size to allow for RPC headers.
- b. the transport can accept packets of at least a minimum size, of approximately 1400 bytes on an Ethernet LAN.

The NFS daemon (nfsd) does not listen on a specified address over transports other than UDP. Therefore, if the nfsd's are killed and restarted it is likely that the new nfsd will listen on an entirely different address causing any previously existing mounts to fail.

The following are options to nfsd that are not documented on the nfsd manual page:

- a start nfsd's over all available connectionless transports
- p *protocol* start nfsd's over the specified protocol
- t *device* start nfsd's for the transport specified by the given device

If the NETPATH environment variable is set, then at least one of the transports in the NETPATH path should be connectionless.

4.2.6 share Resources Command

If a resource is shared with a ro= list and a root= list, any host that is on the root= list will be given only read-only access, regardless of whether that host is specified in the ro= list, unless rw is declared as the default, or the host is mentioned in a rw= list. The same is true if the resource is shared with ro as the default. For example, the following share commands will give read-only permissions to hostb:

```
share -F nfs -oro=hosta,root=hostb /var
```

```
share -F nfs -oro,root=hostb /var
```

While the following will give read/write permissions to hostb:

```
share -F nfs -oro=hosta,rw=hostb,root=hostb /var
```

```
share -F nfs -oroot=hostb /var
```

4.3 RFS

Remote File Sharing (RFS) has been implemented as a file system type under VFS. Even though the new implementation fully supports the pre-Release 4.0 protocol, the new protocol has implications for compatibility and interoperability with previous versions of RFS.

Using `df` with either the `-n` or `-g` options on a remote resource advertised from a pre-SVR4 system will give `unknown` for the `fstype`.

The `-c` option to the `rfs` specific `mount` command function is done by `-o nocaching`.

The `-d` option for the `mount` command is provided for compatibility, but its function is replaced by `-F rfs`, so the `-d` will be removed in a future release.

RFS does not support `mmap`ping of character devices or the allocation of remote swap files.

Users and administrators of RFS clients and servers are reminded that the interpretation of absolute symbolic links encountered on the server can lead to unexpected results because they are relative to the root directory of the client.

Setting up RFS via `sysadm` menus is not currently possible. Menu function for RFS is unpredictable and problematic.

4.3.1 Compatibility with older RFS servers

Facilities new to System V Release 4.0 cannot be provided by an older RFS server. Although an SVR4 RFS client can create a dynamic shared library on a pre-SVR4 server, the shared library cannot be `mmap`ped from that server, because the pre-SVR4 protocol does not support file mapping or paging. The `rename` system call (used by the `mv` command) is not fully supported between SVR4 clients and SVR3 servers. An attempt to `rename` a directory on an SVR3 server will fail with the

error EISDIR.

4.3.2 Swap Control

Swap will not work over RFS (i.e., a swap device or file cannot be created on a resource mounted via RFS). If this is attempted, an ENOSYS error will result.

4.4 RFS: sysadm Initial Remote File Sharing Set Up

It is recommended that you not use the sysadm interface to initiate RFS. Please refer to the chapter on RFS in the *Tips and Troubleshooting* manual.

4.5 Secure RPC

4.5.1 Secure RPC with RFS

The Secure RPC administrative files /etc/masters and /etc/slaves contain the unames of RPC masters and slaves respectively as documented in *Programmer's Guide: Networking Interfaces*. Note, however, that if RFS is used to share the files from the masters or the slaves, then those entries in either file should contain the RFS domain name for that master/server followed by a dot (.) and the uname, i.e., rfsdomain.uname.

One of the Secure RPC administrative tasks of a slave server is to share its /etc directory, writable to its master and readable to its clients.

Note, when NFS is used to share the /etc directory, the share command should be in the form:

```
share -F nfs -r rw=p,root=p,ro=i:j:k /etc
```

where p is the slave's master and i,j, and k are the slave's clients.

4.6 TI_GETINFO

The value of TI_GETINFO and other related ioctl's in sys/timod.h has been changed in SVR4 from TIMOD100 to TIMOD140. These changes should not affect the TLI applications because the SVR4 TLI shared library (libnsl_s) has been re-built using new values. However, if there were private archive versions of libnsl these ioctl changes will affect compatibility.

4.7 sysadm - Remove System Name function

Due to some data parsing problems, the sysadm menu item for removing a system name from a basic networking database does not work. Administrators who want to remove a system from the list of systems that the machine communicates with must

edit and delete the appropriate system name line in the /etc/uucp/Systems file.

4.8 OA&M Modify Menu

Under network_service/selection, the modify menu is not functional. If selected, an error message will appear.

Use the shell level commands to modify the network service, or remove it and then re-add it via OA&M.

4.9 Cable Disconnections During I/O via tty ports

If a cable is disconnected while reading from or writing to an asynchronous communications port (COM1 or COM2), the process will return 0 (zero) bytes in the case of reading (for end of file) and an I/O error in the case of writing.

When the cable is reconnected to the port, the user will continue to get the above return values from read or write, rather than the input or output starting back up again successfully. The only way around this is to close and then reopen the device.

An explanation for this behavior is that upon detecting "*carrier loss*" (disconnecting the cable), the driver sends the message M_HANGUP to the stream head. When the stream head recognizes this, it sets a flag such that from then on, any further reads or writes will return an I/O error or a zero, respectively, to the user. When the carrier comes back (by connecting the cable), the flag is not reset since the streamhead does not know anything about it.

4.10 Configuring Multiple Ethernet Boards

The following describes the procedures for configuring a driver for multiple Ethernet boards of the same type. Multiple Ethernet boards are needed in cases where a UNIX workstation would be used as a gateway between two or more networks. This is technically referred to as a "multi-homed" host. Only the changes needed to configure multiple Ethernet controllers will be discussed here. Refer to the *Network User's and Administrator's Guide* for other configuration issues related to setting up a gateway system.

4.10.1 Multiple PC586, ENET586, iMXLAN586, or N10 Ethernet Board Configuration

The driver for these boards can handle, at most, 4 boards. The following changes need to be made to configure more than one PC586, ENET586, iMXLAN586, or N10 Ethernet board. This procedure assumes that the pc586 driver and inet packages

have already been installed.

4.10.1.1 /etc/conf/pack.d/pc586/space.c file changes

An error in the pc586/space.c file needs to be corrected before multiple boards of this type can be configured. Edit the file /etc/conf/pack.d/pc586/space.c and make the following changes:

At line 68, change PC586_CMAJOR_0 to PC586_CMAJOR_1.

At line 86, change PC586_CMAJOR_0 to PC586_CMAJOR_2.

At line 104, change PC586_CMAJOR_0 to PC586_CMAJOR_3.

After the changes have been made, write the file, exit, and then rebuild the kernel with idbuild.

4.10.1.2 /etc/conf/cf.d/mdevice file changes

Edit the file /etc/conf/cf.d/mdevice and find the line that starts with the word pc586. Its location in the file may vary from system to system. Change this line so that it has the following values:

```
pc586    Ioc    iSchM    pc586  0      xx-xx    1    4    -1
```

Note that there are tab characters between each field. Replace the value shown in the 6th field (xx-xx) with the proper values as follows: Find the largest number in this column for the whole file. (Generally, the largest values will be at the end of the file.) It may even be that the pc586 entry has the largest value. In any case, add one to the largest value and replace the first xx with the result. Now add 3 to the last result, and replace the second xx with this value. For example, suppose the largest value in the 6th column before you started was 50. The pc586 line would then look like this after the change was made:

```
pc586    Ioc    iSchM    pc586  0      51-54    1    4    -1
```

After making the changes, write the file and exit.

4.10.1.3 /etc/conf/sdevice.d/pc586 file changes

Edit the file /etc/conf/sdevice.d/pc586 and add an entry for each additional Ethernet board you will be installing. Each additional line will look the same as the initial line in the file except for the 6th, 9th, and 10th column.

The 6th column is the interrupt level. This must match the interrupt level that the board was jumpered with and each board must have a different interrupt level. The interrupt level can not be used by any other board in your system. You can determine which interrupt levels are in use by looking at the 6th column of all files in

the `/etc/conf/sdevice.d` directory. If an interrupt level is duplicated, an error message will be displayed when you attempt to rebuild the operating system.

The 9th and 10th columns indicate the starting and ending memory address of the Ethernet board. The starting memory address is selectable by jumpers on the Ethernet board. The jumper selections on the board and entry in the `sdevice.d/pc586` file must agree. The following values are possible:

CE000	CFFFF	(default used by Western Digital WD-7000 SCSI disk controller)
C8000	CFFFF	(default used by Adaptec AHA-154x SCSI disk controller)
D0000	D7FFF	
D8000	DFFFF	(default used by WD1003, WD1007 & DPT 2012 disk controllers)
F00000	F27FFF	(not avail. on systems with more than 15 Meg of memory)
F40000	F67FFF	(not avail. on systems with more than 15 Meg of memory)
F80000	FA7FFF	(not avail. on systems with more than 15 Meg of memory)

The memory range specified in each entry can not be used by any other board in the system. You can determine which memory ranges are in use by looking at the 9th and 10th column of all files in the `/etc/conf/sdevice.d` directory. As with the interrupt level, if a memory range of one board overlaps another, an error message will be displayed when you attempt to rebuild the operating system.

4.10.1.4 /etc/conf/node.d/pc586 file changes

Delete all entries in this file and enter the following lines using tabs between each field:

```
clone  pc586_0  c:0  pc586
clone  pc586_1  c:1  pc586
clone  pc586_2  c:2  pc586
clone  pc586_3  c:3  pc586
```

4.10.1.5 /etc/inet/strcf file changes

Edit the file `/etc/inet/strcf`. Near the bottom of this file you should find an entry that looks like this:

```
senet ip /dev/pc586_0 /dev/pc586_1 emd 0#i386/pc586
```

This line should be changed to the following:

```
cenet ip /dev/pc586_0 emd 0 #i386/pc586
```

This changes the entry for the first board. Additional boards would increment the value that follows "emd". For example:

```
cenet ip /dev/pc586_1 emd 1 #i386/pc586
cenet ip /dev/pc586_2 emd 2 #i386/pc586
```

4.10.2 Multiple Western Digital WD8003E Ethernet Board Configuration

The WD8003E driver can handle, at most, 4 boards. The following changes need to be made to configure more than one WD8003E Ethernet board. This procedure assumes that the Western Digital driver (wd) and inet packages have already been installed.

4.10.2.1 /etc/conf/pack.d/wd/space.c file changes

Additional information needs to be added to the the wd/space.c file to use multiple Ethernet boards. Edit the file /etc/conf/pack.d/wd/space.c and find the following lines near the beginning of the file,

```
#define WDIRQ0      WD_0_VECT      /* IRQ value */
#define WDBASEPORT0   WD_0_SIOA      /* Base I/O port */
#define WDBASEADDR0   WD_0_SCMA     /* Base shared memory address */
#define WDBOARDSIZE0  (WD_0_ECMA-WD_0_SCMA)+1    /* Board size */
#define WDMAJOR0      WD_CMAJOR_0   /* Board major device number */
```

Replace all of these lines with the following:

```
#ifdef WD_0
#define WDIRQ0      WD_0_VECT      /* IRQ value */
#define WDBASEPORT0   WD_0_SIOA      /* Base I/O port */
#define WDBASEADDR0   WD_0_SCMA     /* Base shared memory address */
#define WDBOARDSIZE0  (WD_0_ECMA-WD_0_SCMA)+1    /* Board size */
#define WDMAJOR0      WD_CMAJOR_0   /* Board major device number */

#endif /* WD_0 */

#ifndef WD_1
#define WDIRQ1      WD_1_VECT      /* IRQ value */
#define WDBASEPORT1   WD_1_SIOA      /* Base I/O port */
#define WDBASEADDR1   WD_1_SCMA     /* Base shared memory address */
#define WDBOARDSIZE1  (WD_1_ECMA-WD_1_SCMA)+1    /* Board size */
#define WDMAJOR1      WD_CMAJOR_1   /* Board major device number */

#endif /* WD_1 */

#ifndef WD_2
#define WDIRQ2      WD_2_VECT      /* IRQ value */
#define WDBASEPORT2   WD_2_SIOA      /* Base I/O port */
#define WDBASEADDR2   WD_2_SCMA     /* Base shared memory address */
```

```

#define WDBOARDSIZE2      (WD_2_ECMA-WD_2_SCMA)+1          /* Board size */
#define WDMAJOR2           WD_CMAJOR_2                     /* Board major device number */
#endif /* WD_2 */

#ifndef WD_3
#define WDIRQ3            WD_3_VECT                      /* IRQ value */
#define WDBASEPORT3        WD_3_SIOA                      /* Base I/O port */
#define WDBASEADDR3        WD_3_SCMA                      /* Base shared memory address */
#define WDBOARDSIZE3       (WD_3_ECMA-WD_3_SCMA)+1          /* Board size */
#define WDMAJOR3           WD_CMAJOR_3                    /* Board major device number */
#endif /* WD_3 */

```

Now find the following lines located near the end of the file.

```

struct wdparam wdparams[NWD] = {
{
    0,                                     /* board index */
    WDIRQ0,                                /* interrupt level */
    WDBASEPORT0,                            /* I/O port for device */
    (caddr_t)WDBASEADDR0,                   /* address of board's memory */
    WDBOARDSIZE0,                           /* memory size */
    0,                                     /* pointer to mapped memory */
    0,                                     /* board type */
    0,                                     /* board present flag */
    0,                                     /* board status */
    0,                                     /* number of streams open */
    WDMAJOR0,                             /* major device number */
    NSTR,                                  /* number of minor devices allowed */
}
}

```

Replace all of those lines with the following:

```

struct wdparam wdparams[NWD] = {
#endif /* WD_0 */
{
    0,                                     /* board index */
    WDIRQ0,                                /* interrupt level */
    WDBASEPORT0,                            /* I/O port for device */
    (caddr_t)WDBASEADDR0,                   /* address of board's memory */
    WDBOARDSIZE0,                           /* memory size */
    0,                                     /* pointer to mapped memory */
    0,                                     /* board type */
    0,                                     /* board present flag */
    0,                                     /* board status */
    0,                                     /* number of streams open */
    WDMAJOR0,                             /* major device number */
    NSTR,                                  /* number of minor devices allowed */
},
}

```

```

#endif
#define WD_1
{
    0,
    WDIRQ1,
    WDBASEPORT1,
    (caddr_t)WDBASEADDR1,
    WDBOARDSIZE1,
    0,
    0,
    0,
    0,
    0,
    0,
    WDMAJOR1,
    NSTR,
},
#endif
#define WD_2
{
    0,
    WDIRQ2,
    WDBASEPORT2,
    (caddr_t)WDBASEADDR2,
    WDBOARDSIZE2,
    0,
    0,
    0,
    0,
    0,
    0,
    WDMAJOR2,
    NSTR,
},
#endif
#define WD_3
{
    0,
    WDIRQ3,
    WDBASEPORT3,
    (caddr_t)WDBASEADDR3,
    WDBOARDSIZE3,
    0,
    0,
    0,
    0,
    0,
    0,
    WDMAJOR3,
    NSTR,
},
/* board index */
/* interrupt level */
/* I/O port for device */
/* address of board's memory */
/* memory size */
/* pointer to mapped memory */
/* board type */
/* board present flag */
/* board status */
/* number of streams open */
/* major device number */
/* number of minor devices allowed */

```

```
#endif  
};
```

After the changes have been made, write the file and exit.

4.10.2.2 /etc/conf/cf.d/mdevice file changes

Edit the file /etc/conf/cf.d/mdevice and find the line that starts with the word wd. Its location in the file may vary from system to system. Change this line so that it has the following values:

```
wd    I    iSHcfM    wd    0    xx-xx    1    32    -1
```

Note that there are tab characters between each field. Replace the value shown in the 6th field (xx-xx) with the proper values as follows: Find the largest number in this column for the whole file. (Generally, the largest values will be at the end of the file.) It may even be that the wd entry has the largest value. In any case, add one to the largest value and replace the first xx with the result. Now add 3 to the last result, and replace the second xx with this value. For example, suppose the largest value in the 6th column before you started was 50. The wd line would then look like this after the change was made:

```
wd    I    iSHcfM    wd    0    51-54    1    32    -1
```

After making the changes, write the file and exit.

4.10.2.3 /etc/conf/sdevice.d/wd file changes

Edit the file /etc/conf/sdevice.d/wd and add an entry for each additional Ethernet board you will be installing. Each additional line will look the same as the initial line in the file except for the last 5 columns (6th, 7th, 8th, 9th, and 10th).

The 6th column is the interrupt level. This must match the interrupt level that the board was jumpered with and each board must have a different interrupt level. The interrupt level can not be used by any other board in your system. You can determine which interrupt levels are in use by looking at the 6th column of all files in the /etc/conf/sdevice.d directory. If an interrupt level is duplicated, an error message will be displayed when you attempt to rebuild the operating system.

The 7th and 8th columns indicate the starting and ending I/O address used by the board. The starting I/O address is selectable by jumpers on the Ethernet board. The jumper selection and entry in the sdevice.d/wd file must agree. The following values are possible:

Start	End	Start	End
200	- 21F *	300	- 31F
220	- 23F	320	- 33F *
240	- 25F	340	- 35F *
260	- 27F	360	- 37F *
280	- 29F	380	- 39F
2A0	- 2BF	3A0	- 3BF
2C0	- 2DF	3C0	- 3DF
2E0	- 2FF *	3E0	- 3FF *

The I/O range can not be used by any other board in the system. The ranges with asterisks by them are values that are often used by other parts of the system and should be avoided. You can determine which memory ranges are in use by looking at the 7th and 8th column of all files in the /etc/conf/sdevice.d directory. As with the interrupt level, if one I/O range overlaps another, an error message will be displayed when you attempt to rebuild the operating system.

The 9th and 10th columns indicate the starting and ending memory address of the Ethernet board. The starting memory address is not selected by jumpers on the Ethernet board but by the values specified in this file. The following values are possible:

C8000	C9FFF	(default used by Adaptec AHA-154x SCSI disk controller)
CA000	CBFFF	
CC000	CDFFF	
CE000	CFFFF	(default used by WD7000 - FASST SCSI disk controller)
D0000	D1FFF	
D2000	D3FFF	
D4000	D5FFF	
D6000	D7FFF	
D8000	DFFFF	(default used by WD1003, WD1007 & DPT 2012 disk controllers)

The memory range can not be used by any other boards in the system. You can determine which memory ranges are in use by looking at the 9th and 10th column of all files in the /etc/conf/sdevice.d directory. As with the interrupt level, if one memory range overlaps another, an error message will be displayed when you attempt to rebuild the operating system.

4.10.2.4 /etc/conf/node.d/wd file changes

Delete all entries in this file and enter the following lines using tabs between each field:

```
clone    wd0    c:0    wd
clone    wd1    c:1    wd
clone    wd2    c:2    wd
clone    wd3    c:3    wd
```

4.10.3 Rebuilding the Operating System

If you will be configuring the system to be a gateway between two or more networks, you should make the changes outlined in the "Network User's and Administrator's Guide" for configuring a gateway now.

After all files have been modified, rebuild the operation system with the following command:

```
/etc/conf/bin/idbuild
```

If there were any conflicts on interrupt level, I/O address range, or memory address range, an error message will be displayed indicating where the problem is. It will be up to you to select new values that don't conflict with other parts of the system.

If there were no errors, shut the system down with the following command:

```
shutdown -g0 -y
```

When the system prompts that you may reboot the system, turn power off and install the additional Ethernet boards you have just configured the system to use.

5. System Administration and Maintenance

5.1 Improved Backup and Restore Operations

5.1.1 SVR3 Compatibility

System administrators must recreate the online backup schedule using the `bkreg` command.

5.1.2 Future Directions

The `ckbupsd` command will not be supported in future releases.

5.2 Backup and Restore

When using the `backup/restore` commands, a specific device name must be specified, not just the device group. If for some reason, the device group also must be specified, the "spool" entry in `/etc/device.tab` must be moved to the end of the file.

Save a copy of the `/etc/device.tab` file as it is recreated each time the system boots.

Labels must be specified for all backups and there must be enough for the number of volumes to be created. The option to overwrite the label, `-o`, should not be used.

5.3 Backup/Restore Work-arounds

5.3.1 `-o` Method Option

The `-o` method option overrides label checking on `backup`. The override allows a backup to occur to a tape or disk that has a label different from that specified for the operation. Because `restore`, and `rsoper` always check labels, non-NUL labelled media should always be supplied in any backup operation, even with label checking override.

5.3.2 `Fimage Archives`

The `fimage` method restores an entire file system. If such an archive is online, and the file system is unmounted, the restore will occur automatically. A partition may be unintentionally overwritten if the administrator requests a `restore` and such an archive is available.

5.3.3 `Backup Strategy`

Both `fimage` and `ffile` methods are complete file system backup strategies and are mutually exclusive. Correct strategy is to use either `fimage` or `ffile` in conjunction with `incfile`.

5.3.4 `Viewing Archive Labels`

The `rsoper -n -d` command allows an administrator to view the label of the archive on the device specified. This is useful if external labels are damaged or a backup history is lost.

5.4 /etc/bkup/rsstatus.tab: Not Created During Installation

The file `/etc/bkup/rsstatus.tab` is not created when the system is installed. If you attempt to query this file using the Extended Restore function under OA&M, you will get a message to this effect. This is not serious as the file will be created when the first Extended Restore is executed under OA&M.

5.5 Console Logging

The `/dev/console` and `/dev/syscon` special files are still available for use in this release. Access to a console by `/dev/console` will be unavailable in a future release.

5.6 System Administration Menus (sysadm)

A new screen-based administrative interface, `sysadm`, is introduced in SVR4. This interface uses FMLI utilities and can be operated with the function keys on most terminals/keyboards.

5.6.1 Compatibility

The administrative interface, `sysadm`, is dependent upon the FMLI package. Tasks may be accessed directly from the shell level express mode.

5.6.1.1 Function Key Resets

In a limited number of cases, users may find that keyboard function keys may need to be reset after the termination of a `sysadm` session.

5.6.2 Future Directions

The administrative logins (i.e. `sysadm`, `setup`, `powerdown`, `makefsys`, `mountfsys`, `umountfsys`, and `checkfsys`) will not be supported in the next release. Commands with the same names, however, will be supported.

All support for the pre-SVR4 menu items and pre-SVR4 express mode names will be removed in a future release.

5.7 Configuring a Second Serial Port

The default configuration is to have the second serial port disabled. First, you need to add the node in `/etc/conf/node.d/asy` by coping the entry for `tty00` and changing the `00` to `01`. The file looks like this before:

```
asy  tty00      c  0
asy  term/00    c  0
asy  tty00s    c  0
asy  tty00h    c  128
asy  term/00s   c  0
asy  term/00h   c  128
```

And this after the change:

```
asy  tty00      c  0
asy  term/00    c  0
asy  tty00s    c  0
asy  tty00h    c  128
asy  term/00s   c  0
asy  term/00h   c  128
asy  tty01      c  0
asy  term/01    c  0
asy  tty01s    c  0
asy  tty01h    c  129
asy  term/01s   c  0
asy  term/01h   c  129
```

To enable it as a tty port using the `sysadm` interface follow the steps below:

1. In the `ports/port_monitors/add` form enter the following

Port monitor tag: `portmon`
Port monitor type: `ttymon`
Command to start the port monitor:
 `/usr/lib/saf/ttymon`

Then save (F3).

2. Go to the `ports/port_monitors/enable` menu and enable `portmon`.
3. Go to the `ports/port_services/add/add_to_one/portmon` form and enter the following:

Service tag: portmon
Service invocation identity: root
Port/service state: ENABLED
utmp entry to be created for this service? Yes

Then save (F3).

4. On page 2 of the same form, enter /dev/tty01 for the name of the tty device. Then save (F3).
5. Now connect a terminal to the COM2. If a terminal is connected before this step, you will not get a login prompt.

5.8 User and Password Administration

The pwck command knows only about /etc/passwd, and not /etc/shadow.

The pwconv command will not be updated in future releases to work with new identification and authorization database files.

The passmgmt command will be removed from distribution in the next release. The functionality has been replaced by the SVR4 useradd, usermod, and userdel commands.

5.9 Video Display Card

5.9.1 No Integral Video Card

On systems with no integral video card (i.e. only terminals connected to the serial ports or X-terminals connected via Ethernet), you will need to edit the file /etc/ap/chan.ap to remove the line for the keyboard driver (major device number 5). If no integral display is present, the system will boot, but no output will be displayed.

5.9.2 EVGA Mode Changes

For the video cards that are described on the evgainit manual page, there are new mode switching constants in kd.h. They are all prefixed with SW_GEN_ followed by the resolution. These should be used to set any of these video cards to graphics modes with resolutions higher than 640x480. For example, to set the video card to 800x600, the SW_GEN_800x600 constant would be used. The ioctl to set the mode will succeed if the video card in question supports the specified resolution, otherwise it will fail with errno set to ENXIO. In order to set the video card mode,

first open `/dev/video` and then do an `ioctl` on the file descriptor with the appropriate mode switching constant. For example:

```
fd = open("/dev/video", O_RDWR);  
  
ioctl(fd, SW_GEN_800x600, 0);
```

5.10 Quick Terminal Set-up and vt Devices

The OA&M interface and Quick Terminal Set up is used to quickly spawn a `ttymon` or `getty` service on a particular port, but quick terminal set up does not look for `vt` devices. `vtlmgmgr` should be used.

5.11 pkgadd

There is a problem with the OA&M `pkgadd` if the files contained in the OA&M package have group ownership values not contained in the system's `/etc/group` file. Enter the group owner ID values in the `/etc/group` file.

When using `pkgmk` to create an installable package, the capacity specified should be slightly less than the true capacity of the device. The simplest way to do this is to add `-1 small size` to the `pkgmk` command line. Where *small size* is roughly equal to the size of the `pkginfo` file (in 512 byte blocks) plus 1. If the new value is still too large, a warning about no space being left on the device will be sent to the console.

If a package removal fails, attempt to remove the package again. If the removal still fails the second time, then try to remove the lock file as outlined here. To remove the lock file do the following:

1. `cd /var/sadm/pkg`
2. `cd pkgnname`, where *pkgnname* is the directory corresponding to the package you are trying to install. Check the `NAME` field in the `pkginfo` file to assure the correct directory has been selected
3. Remove the file called `!R-Lock!`.

5.12 No Help for Some OA&M Screens

The following menus have no Help text:

- `restore_services/extended/status/full`

6. Real Time Processing

6.1 High-Resolution Timers

6.1.1 Source Compatibility

In SVR4, `gettimeofday` and `settimeofday` take one argument. The second argument used with the respective BSD calls (and in the BSD compatibility library) has long been obsolete and is ignored.

In BSD, the `alarm` and `sleep` functions are written using `setitimer` [see `getitimer(3C)`]. In SVR4, `alarm` is independent of `setitimer`. In SVR4, a `sleep` following a `setitimer` wipes out knowledge of the user signal handler; do not use `setitimer` with `sleep`.

6.1.2 Future Directions

These routines will be included in SVR4 until the POSIX P1003.4 standard on Real Time System Interfaces is completed.

7. C Programming Language

7.1 External Data/Automatics Order

As always, the order of external data and automatics on the stack should not be relied on. In particular, the compilation system is free to do register allocation.

7.2 Floating Point Arithmetic

All floating point operations are performed in double extended precision until explicitly stored in memory. The following example illustrates the problem; it will print `not equal`.

```
#include <math.h>

main()
{
    double d = (nextafter(1.0, 1.1) + 1.0) / 2.0;

    if (d != ((nextafter(1.0, 1.1) + 1.0) / 2.0))
        printf("not equal\n"); /* expression is between [1.
    else
        printf("equal\n");
}
```

- `backup_services/extended/status/full`
- `applications`
- `network_services`

Consult *Network User's and Administrator's Guide* for help.

5.13 mkpart Command

Due to problems in the `mkpart` command, users will be required to use the new disk administration commands provided with SVR4. The current use of `mkpart` is limited and consequently most users will not be affected. If you require the use of a `mkpart` option however, there are some alternative commands.

- For displaying VTOC, pdinfo, or alternates tables, use the `prtvtoc` command.
- For rewriting the boot code to the hard disk, use the `disksetup` command with the `-b` option.
- For modifying the contents of the VTOC, use the `prtvtoc` command with the `-f` option (to generate the VTOC file) and the `edvtoc` command to modify the VTOC.

Manual page for `prtvtoc` is in *Tips and Troubleshooting* manual.

5.14 Adding a Printer with sysadm

After adding a printer through `sysadm` you get a message that a frame cannot be found nor closed. If you respond `no` to the default baud rate and parity, and complete the `add`, you receive a "*cannot close message*" for the following frame:

```
/usr/sadm/sysadm/add-on/lp/printer/printer/add/Text.confir
```

Please note that the printer is successfully added.

5.15 Rebooting After Daylight Saving Time Loses Adjustment

If your machine is up during the time adjustment for Daylight Savings Time, the `date` command shows you the correct local time (the clock is adjusted properly one hour ahead for Daylight Savings). If the system is brought down normally via `shutdown` and rebooted, the time adjustment for Daylight Savings is lost. The workaround is to make sure you reset the time and date with either the `date` command or `sysadm` after Daylight Savings Time.

7.3 ifiles No Longer Supported

The `ifiles` feature of the link editor command language is no longer supported. You should switch to the new `mapfiles` feature, which subsumes much of the functionality.

7.4 Optimizer and asms

The optimizer does not handle `asms` that change the depth of the stack.

7.5 Performance Tradeoffs

The initialized arrays `_job` and `_ctype` are defined in the dynamic C library since they are referenced by many of the library functions. They may also be referenced directly from the user's code through macros such as `getchar` and `isdigit`. Since the user's code is not typically compiled as position independent, space for these symbols must be allocated in the `a.out`'s data segment. At process startup time, the dynamic linker changes the global offset table entries in the library to point to the `a.out`'s symbols, if present, so that both the library and the `a.out` will reference the same object.

`_job` and `_ctype` are fairly large arrays, and the method chosen to initialize the arrays in the `a.out`'s data segment can have an impact on overall system performance. In one method, both `libc.so` and `libc.so.1` are built with the files (`data.o` and `_ctype.o`) containing the initialized data. Executable files linked with `libc.so` will have the initialized arrays in their `.data` sections.

The alternative method is to replace the definitions in the archive `libc.so` with uninitialized objects (`.bss` symbols), and add code to the dynamic linker to copy the initialized data from the shared library into the `a.out`'s `.bss` section at process startup. This alternative typically makes executable files smaller, and trades off disk I/O for longer startup times.

7.6 Shared Libraries

Existing static shared libraries will continue to work with this release. You cannot, however, create new ones, that is, the `mkshlib` and `chkshlib` commands are no longer supported. This functionality has been completely replaced with the new dynamic shared library feature.

7.7 Commands and Functions

7.7.1 ctrace Command

The **ctrace** command does not handle **asms**.

7.7.2 dump Command

The **dump** command returns a status code indicating successful completion (return code set to zero) despite the following cases of bad input:

- number out of range
- invalid range
- file name not found
- no such file or directory
- invalid file type
- bad line info section

7.7.3 ld Command

- Error messages from the link editor that refer to I/O errors may be caused by an inability to create a file. This may be caused by such things as file space limitation, permission problems, or **ulimit** problems.
- Error messages from the link editor that specify "vm stats" errors or "output file space" may result from running out of system swap space.

7.7.4 lint Command

- **lint -p** gives the diagnostic "pointer cast may be troublesome" when two pointers differ only by a **const**. (For example, **const int *** versus **int ***.)
- **lint's printf** format checks do not recognize positional parameters. Therefore, **printf("%1\$s", s);** will yield the warning "too many arguments for format."

7.7.5 lprof Command

- **lprof** does not ignore C code included via a header file. When it encounters a function that is defined in a header file, it begins outputting line numbers and line counts. Since this code does not appear in the original C file, the line counts will be off. A workaround is to move the C functions from the header file(s) into **.c** files and compile them as separate modules.

- lprof will fail if the executable being profiled is built using two or more object files with the same basename, e.g., directory1/fun.o and directory2/fun.o. It will not be able to read the .cnt file created when the program is executed.

7.7.6 nm Command

The nm command returns a status code indicating successful completion (return code set to zero) when it encounters and warns about a truncated file.

7.7.7 SCCS Commands

Using the delta command on a text file that contains a line whose line length exceeds 1024 characters will cause an infinite loop.

7.7.8 sdb Command

- Assembler routines that use the frame pointer and argument pointer registers in non-standard ways may cause erroneous stack traces in sdb. The __doprnt function is an example.
- Header files that include executable code may cause confusion with various tools (such as sdb) about line numbers. The result is that output from the tools relating to a function described in a header file may be associated with the wrong line number.
- When a process is grabbed (via sdb /proc/123, for example), you may examine variables, instruction step, quit, continue, or kill the process. However, breakpoints and some statement steps will not work. The process will be killed if it hits a breakpoint, either one you set explicitly or one set by sdb to statement step over a function call.

7.7.9 mmap Function

Dynamically linked a.outs that depend on a shared library that resides on a file system that does not support the mmap function call (e.g., the BFS file system) will not execute.

8. Line Printer Spooling Utilities

8.1 Networking

The new Line Printer Spooling Utilities networking feature does work, although it is possible that stressing the network capability may result in suspending LP network

activity requiring the administrator to restart lpsched.

It is also possible that the networking feature may fault when attempting to network with a BSD system.

If a user's remote print request is cancelled by the administrator on the remote machine, the user is not notified by mail. Cancelling the request from the local machine generates the correct mail notification. In both cases the request is successfully cancelled.

8.2 BSD Compatibility Commands

The BSD compatibility commands lpc, lpq, lpr, and lprm are operational, but may not be fully functional. lprm has only partial functionality. Please use the Line Printer Spooling commands instead.

8.3 Line Printer Spooling Utilities over RFS

Due to directory structure changes, SVR4 LP is not compatible with SVR3 LP across RFS. Note that when you run the Line Printer Spooling Utilities over an RFS network, all machines in the network must be running the same version of LP.

8.4 lpstat Command

If the output of lpstat -o -l always shows that a particular lp request is cancelled, then the following script should be run as root:

```
lpshut
cd /var/spool/lp
find requests tmp ! -type d ! -name .SEQF -exec rm {} \;
/usr/lib/lp/lpsched
```

9. BSD Compatibility Notes

The BSD Compatibility package supports only 4.3BSD user level applications. It does not include machine and architecture specific files (e.g. access routines to mtab and fstab are not supported). Header files and functions for BSD 4.1/4.2 compatibility are also not supported. A list of unsupported features are as follows:

1. The utmp structure in BSD and SVR4 are incompatible.
2. The flock() routine is not supported along with its parameters: LOCK_EX, LOCK_NB, and LOCK_UN. It is recommended that the lockf() routines be used instead.

3. `_IOSTRG` defined in BSD `stdio.h` is not supported.
4. `ranlib` and `<ranlib.h>` are not supported due to the new format of libraries in SVR4.
5. `IFDIR`, `IFMT`, `IFREG`, etc are defined in `sys/inode.h` in 4.3BSD, in SVR4 these flags are defined in `sys/fs/s5inode.h` and `sys/fs/ufs_inode.h`. Since SVR4 supports more than one file system, these values cannot be provided in `sys/inode.h`.

10. X Windows Notes

10.1 Initial Setup

If you do not explicitly set the `XDISPLAY` environment variable in your startup script (e.g. `.profile`), then you will need to:

1. Run `/usr/X/adm/oladduser`
2. `... .profile`

This will set you `XDISPLAY` variable on login along with putting `/usr/X/bin` in your path.

10.2 Set Up Instructions

`xterm` has OPEN LOOK³ features; cut-and-paste is handled differently than in previous releases. Please read the *OPEN LOOK User's Guide* carefully for details.

In particular, the user may notice that `CTRL-c`, `CTRL-v`, `CTRL-p` and other keystrokes are not interpreted as expected. `xterm`, as a part of its OPEN LOOK Look and Feel, interprets these keystrokes as shortcuts for some OPEN LOOK commands. To get around this (which `vi` users may desire, and `emacs` users have to do), change the properties either using `olwsm` (WARNING: this program WILL alter your `.Xdefaults` file) or change the `.Xdefaults` file manually. Adding the lines:

3. OPEN LOOK is a trademark of AT&T

```
*copyKey:      Mod1<c>
*cutKey:       Mod1<x>
*pasteKey:     Mod1<v>
*propertiesKey: Mod1<p>
*stopKey:      Mod1<s>
*undoKey:      Mod1<u>
```

to your `.Xdefaults` will map the related keystrokes from their control key combination to a counterpart using the ALT key.

For those of you who have old X clients, you need to watch out for one major change from X11R2 to X11R3—fonts. The locations for accessing fonts has changed, and the actual fonts available may be different than before. If an old client dies with an X error, look carefully at the error message to determine if the problem is just that a previously available font is either no longer available, or available via another name.

Experienced users with their own `xinitrc` files should make sure that the file contains one client that will "stay around" for the duration of the session. Do NOT have all clients put in the background.

10.3 xterm Command

10.3.1 xterm does not always startup iconically

xterm with the `-i` option does not always startup the xterm iconically.

10.3.2 Passwd command does not work in xterm

If you type `passwd` while in an `xterm` you get the following message back from the system.

```
Usage:
passwd [-s] [name]
```

You can not change your password on an X window unless you type:

```
passwd your-login-name
```

To work around this problem, you can switch to the console via a vt switch <ALT-SYSREQ> H, then execute the `passwd` command. To get back to the `xterm`, type <ALT-SYSREQ> P.

10.4 Olfm Does Not Change Directory

After you have searched into a sub-directory and then perform a menu entry from a file selection, the command that is performed from the `.olfmrc` file is not done in the current directory. The current working directory remains in the directory that `olfm` was started in.

10.5 Known Problems

1. The client `xwd` has been known to fail on occasion when a window is selected for "dumping". If this happens, the user may wish to use either `xwininfo` or `xlswins` to determine the numeric window id for the desired window; this numeric window id may then be passed to `xwd` using the `-id` option or use `olprintscreen`.
2. Old (pre-X11 Release 3) clients often worked in prior releases of X despite failing to follow some explicit requirements of the X protocol. If you run an old client, and see an error message like "*X Protocol error detected by server: integer parameter out of range...*", then it is likely that the client contains an instance of the bug. The client should still be usable if the server is instead invoked with the "backward (or bug) compatibility" option. This may be done by starting `xinit` with the `bc` option, thusly,

```
xinit -- bc
```

3. Users should be sure to provide some `sleep` time between invocations of successive clients in their `.xinitrc` or `.olinitrc` file. Attempting to start many clients at once could result in clients never showing up on the display.
4. `xpic`⁴ will report that no fonts are available.

To have access to fonts, place the following lines in a file called ".xpic" in the user's home directory:

4. Copyright University of Toronto 1988, 1989. Written by Mark Moraes

default Times-Roman 10
Times-Roman *-times-medium-r-normal-*
Times-Italic *-times-medium-i-normal-*
Times-Bold *-times-bold-r-normal-*
Times-BoldItalic *-times-bold-i-normal-*
Courier *-courier-medium-r-normal-*
Courier-Bold *-courier-bold-r-normal-*
Courier-BoldOblique *-courier-bold-o-normal-*
Courier-Oblique *-courier-medium-o-normal-*
Helvetica *-helvetica-medium-r-normal-*
Helvetica-Bold *-helvetica-bold-r-normal-*
Helvetica-BoldOblique *-helvetica-bold-r-normal-*
Helvetica-Oblique *-helvetica-medium-r-normal-*
Symbol *-symbol-medium-r-normal-*
Special *-symbol-medium-r-normal-*
Roman *-times-medium-r-normal-*
Bold *-times-bold-r-normal-*
BigBold *-times-bold-r-normal-*
Image-Roman *-times-medium-r-normal-*
Image-Italic *-times-medium-i-normal-*

5. *xpic* is known to have problems deleting splines and some lines.
6. There is an odd interaction between some fonts and *xterm* such that extraneous bits can be left at the bottom of a line. The user may not even notice the problem, depending on the fonts used and the resolution of the screen.
7. *xmj*⁵ installs its own colormap. This can be distracting when entering and leaving the game's window. No harm should result from this, however.
8. The game *roids*⁶ starts up with a very small window, and it is too slow to be useful. The user may wish to remove it.

5. Original source Copyright 1988 Exceptions
X11 portions Copyright 1989 Concurrent Computer Corporation
6. Copyright 1989 Digital Equipment Corporation

9. *xcalendar*⁷ does not know the correct location for its help file. To work around the problem, modify the application defaults file (*/usr/X/lib/app-defaults/XCalendar*) to reference the correct help file. Add the line:

*XCalendar*helpFile: /usr/X/lib/xcalendar.hlp*

10. Some DOS applications have been known to set the mouse in an odd state (whether run under a DOS emulator, or run under DOS before booting off the UNIX partition). If the mouse behaves incorrectly when X is started, disconnect the mouse, then reconnect it. This should solve the problem.
11. Can't cut and paste between the new *xterm* and *X11R2* clients. To workaround, paste into an *xedit(1)* window and save the text to a file, then display the file in an *xterm* and cut from there.

7. Copyright 1988 Massachusetts Institute of Technology

